

# ThermoLoop

---

*Modular, Closed-Loop Thermal Recycling  
for AI Data Centers*

Cade Kaminski · Sam Wolpert



# Problem

49 GW

US generation  
shortfall by 2028

120 kW

per AI rack currently  
(vs. 5-10 kW legacy)

300K

gallons water/day  
per 100 MW facility

33%

retail electricity  
price increase since 2019

## Constraints

Power, cooling, water, waste heat, carbon, community are addressed in silos, creating compounding inefficiencies. They share one root cause:

**Thermal energy is treated as waste instead of a resource**

# Why Not Other Big Ideas?

The digital infrastructure industry has proposed bold solutions to the AI power crisis. Many are promising long-term, but none can be deployed fast enough to win the AI race today.

## Small Modular Reactors

Billions invested, but SMRs are 5+ years from US commercial operation. NRC licensing takes years. No data center on Earth is powered by an SMR today.

## Underwater Data Centers

Microsoft's Project Natick was cancelled in 2024. Sealed underwater hardware upgrades are neither pragmatic nor economical in a world where GPU architectures change every 18 months.

## Data Centers in Space

Many new startup companies, but launch costs need to drop 10x from today's rates before it is economical. Latency, maintenance, and bandwidth remain unsolved at scale.

***ThermoLoop uses technology that exists today, deploys in 12 to 14 months, and solves six constraints at once.***

# The Solution: ThermoLoop

A four-stage, closed-loop thermal recycling architecture that captures waste heat from AI server racks and converts it into the facility's own cooling supply or supplemental electricity.

## Stage 1

### Heat Capture

Two-phase D2C cold plates absorb GPU heat. Coolant exits at 50-65°C.

## Stage 2

### Heat Recycling

Adsorption chiller or ORC turbine converts waste heat to cooling or power.

## Stage 3

### AI/Ignition Dispatch

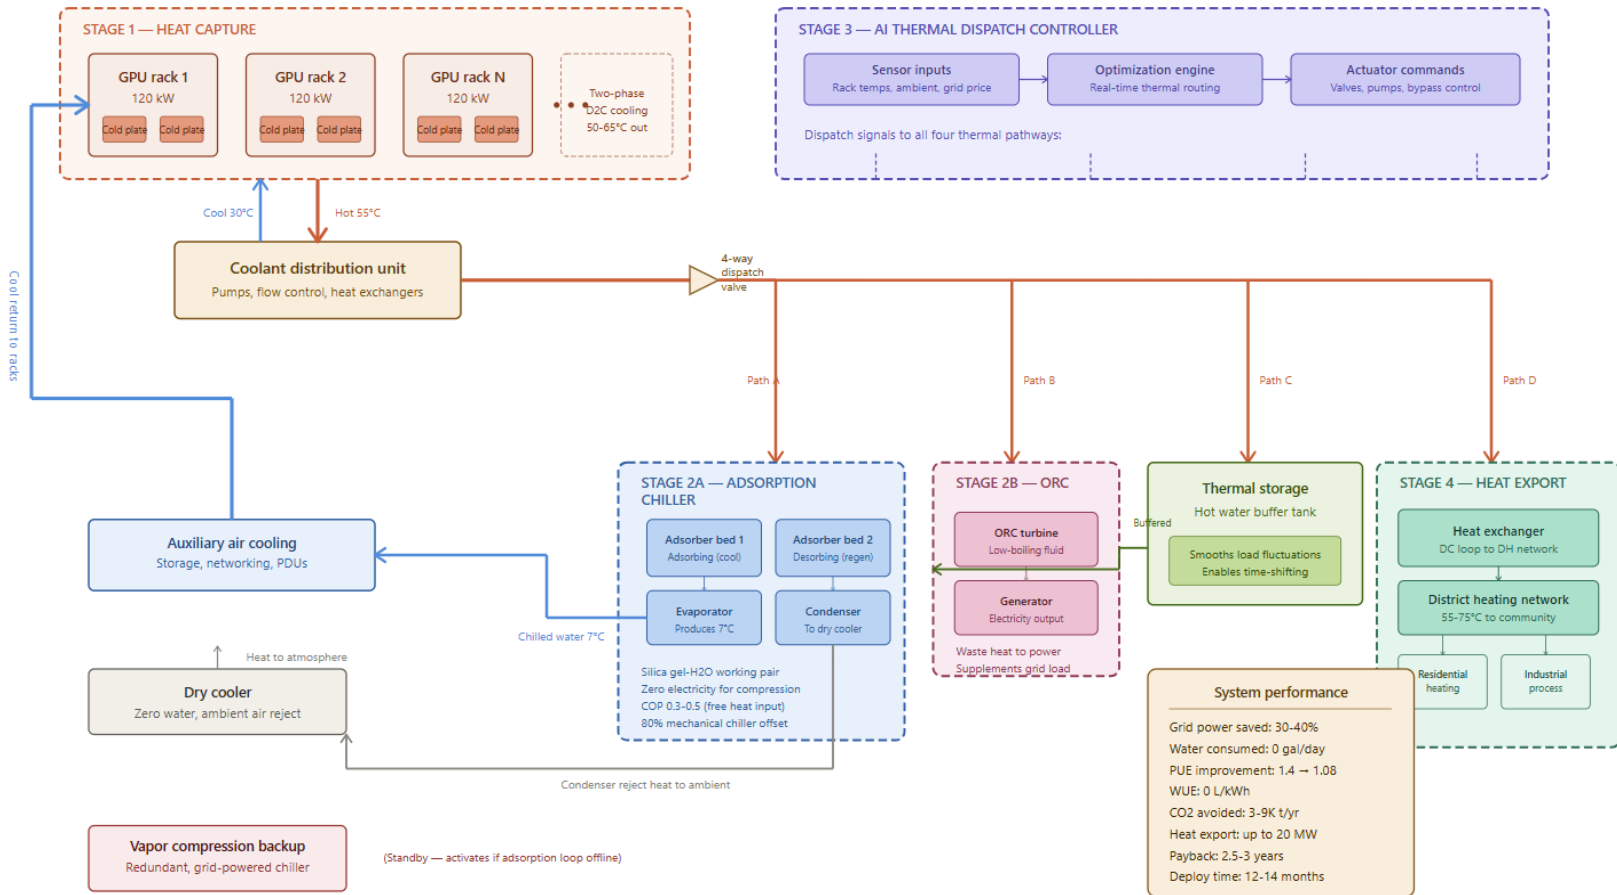
Software routes heat in real time to highest-value destination.

## Stage 4

### Heat Export

Surplus heat goes to district heating, greenhouses, or industry.

***Every component exists commercially today with no experimental technology required***



**LEGEND:** — Hot coolant (50-65°C) — Chilled / cool return (7-30°C) — Condenser reject — Control signal — Stage boundary

All stages modular and skid-mounted. Dispatch selects adsorption chiller (cooling priority) or ORC (power priority) in real time. Backup VCC ensures uptime if thermal loop is offline.

# Stage 2: Heat Recycling (The Core Innovation)

## Path A: Adsorption Chiller

*When cooling is the priority*

- Uses silica gel-H<sub>2</sub>O solid sorbent material
- Captures 50–65°C low-grade waste heat from server cooling loops
- Produces chilled water at ~7°C for auxiliary cooling
- Zero electricity consumed for refrigeration cycle
- Reduces mechanical cooling load by up to 80.5%, directly cutting the facility's largest electricity draw
- COP 0.3-0.5 (input heat is free, so COP is less critical)

## Path B: Organic Rankine Cycle

*When power recovery is the priority*

- Uses low-boiling-point organic working fluid
- Thermal efficiency of 8–15% converting waste heat to electricity
- Supplements grid load, reducing net facility draw
- Ideal when electrical capacity is the binding constraint
- AI dispatch selects ORC when electricity prices peak

# Impact at a Glance

Challenge	ThermoLoop Response	Quantified Benefit
Power / Grid	Cooling load shifts to recycled waste heat; ORC returns supplemental power	30-40% less grid power per 100 MW facility
High-Density Cooling	Two-phase D2C for 100+ kW racks; adsorption chiller for auxiliary	Up to 82% less chip cooling energy; 80% chiller offset
Water	Fully closed-loop; zero evaporative consumption	~33M gal saved/yr; WUE approaches 0
Waste Heat	Waste heat drives adsorption cooling or ORC generation	70-90% of electricity recoverable as thermal energy
Carbon	Lower grid draw + exported heat displaces fossil heating	3,000-9,000+ tonnes CO2 avoided/yr
Community	Heat-as-a-service revenue; zero water/grid burden	10,000+ homes heated per 20 MW IT load
Network Availability	Modular skid-mounted design for frontier markets	12-14 month deploy vs. 4-10 yr grid queue

# Cost Analysis

**\$3.6M**

Total CAPEX  
(20 MW facility)

**\$9.5M/yr**

Total Annual  
Benefit


**< 1 yr**

Simple  
Payback

**\$91M**

10-Year Net  
Cash Flow

**Economic  
Analysis Spreadsheet**



## Annual Benefit Breakdown

GPU capacity reclaimed: \$4,900K  
Heat export revenue: \$2,628K  
Electricity savings: \$1,502K  
Carbon credits: \$429K

# Risks & Opportunities

## RISKS

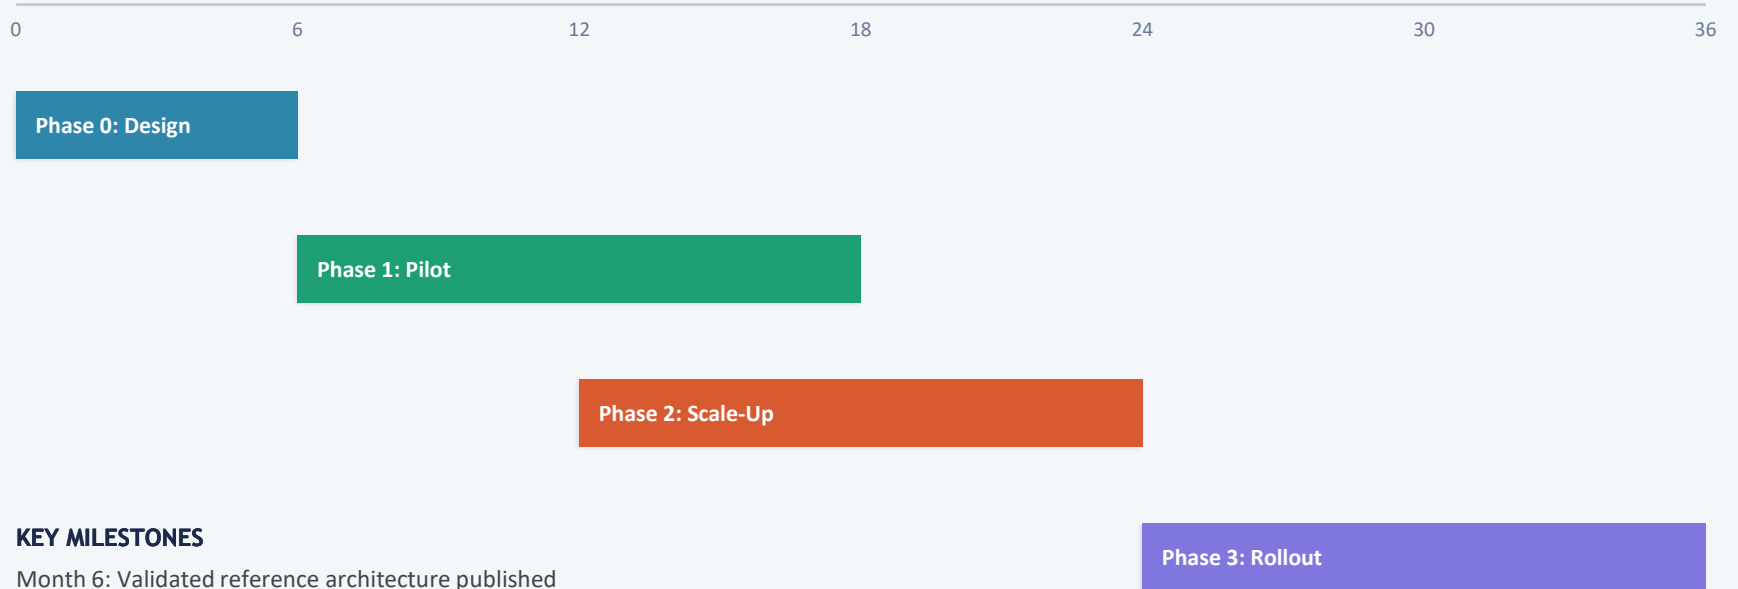
- Adsorption COP lower than vapor compression, but input heat is free, so marginal cost is near zero
- CAPEX premium over conventional cooling; mitigated by 2.5 - 3 year payback and reclaimed GPU capacity value
- Multi-loop system adds operational complexity; mitigated by AI dispatch automation and modular skid packaging
- District heating demand varies seasonally; Stage 4 is optional and thermal storage enables time-shifting

## OPPORTUNITIES

- Germany mandates 10% waste heat reuse by July 2026; US regulation will surely follow
- Heat-as-a-service creates recurring revenue from what was previously waste, offsetting cooling OPEX entirely
- 64% of US data center capacity under construction is in frontier markets, ideal for ThermoLoop integration
- Liquid cooling market growing to \$7B by 2029. ThermoLoop rides existing adoption curves.

*Mitigation: Phased rollout de-risks each stage. VCC backup chiller ensures uptime if adsorption loop goes offline.*

# Implementation Timeline



## KEY MILESTONES

- Month 6: Validated reference architecture published
- Month 18: 30%+ cooling energy reduction demonstrated at pilot
- Month 24: First commercial heat-as-a-service contract signed
- Month 36: Multi-site deployment; industry standard proposal submitted

# Why ThermoLoop Wins

## Systems Innovation

Not a component invention. Two-phase cooling, adsorption chillers, and ORC turbines all exist. ThermoLoop is the unified architecture that makes them work together.

## Economic Necessity

30-40% more GPUs per megawatt without drawing additional grid power. The strongest argument is economic, not environmental.

## Deployable Today

Skid-mounted, factory-tested modules. 12-14 months from decision to operation. No SMRs, no fusion, no regulatory mandates required.

**The company that deploys 30% more GPUs per megawatt wins the AI race.  
ThermoLoop is designed to be that system.**

# Questions?

Cade Kaminski · Sam Wolpert